Information Immune Systems

Dennis L. Chao and Stephanie Forrest
Department of Computer Science
University of New Mexico
Albuquerque, NM 87131 USA
{dlchao,forrest}@cs.unm.edu

Abstract

Many people are exposed to more information than they can process effectively. We describe an approach to building an *information immune system* that eliminates undesirable information before it reaches the user. This approach is inspired by natural immune systems that protect us from pathogens. The potential applications of an information immune system include filtering out undesirable data, generating a variety of solutions to a design problem, and finding consensus solutions to problems.

1 Introduction

Information overload is inevitable in a world that produces over an exabyte (one billion gigabytes) of information per year [28]. We will continue to produce and consume more and more information, so we must find innovative ways to manage it. Although finding and managing information are active research areas, much of the effort is directed towards active strategies such as information retrieval. Although these techniques help individuals locate desirable information, they also accelerate the information glut.

In this paper, we outline the features of an *information immune system* (IIS)[1] that could help people deal with the glut of data. We draw inspiration from natural immune systems that protect us from a seemingly limitless number of possible invaders such as bacteria, viruses, and parasites. We believe that an IIS can be constructed to eliminate undesired information after detecting it in a manner analogous to the natural immune system's. Such an IIS would be situated between an individual and a stream of information as a mediator. Instead of actively bringing more pieces of information to our attention, it will quietly censor unwanted data.

[1]The term "information immune system" was introduced by Neil Postman (in [39] and [40, page 63]).

An IIS should be capable of learning what kinds of information a user wants and discarding the rest. The task of distinguishing what is desirable is a difficult one. We propose taking one of the approaches used by our natural immune systems, which can "remember" a pathogen that infects us so it can eliminate it more quickly in future encounters. An IIS can do this by storing examples of rejected information and censoring similar data. If the memory of the system is too specific, this approach is likely to be ineffective. Pathogens and information can mutate over time, and our immune systems must be able to generalize. Therefore, both the natural and the information immune systems must also learn to eliminate *related* pathogens while taking care not to harm anything else.

An extension to a personal IIS is a group IIS. If one uses the IISs of many individuals in serial to filter a stream of information, the only information that can survive all IISs is the information that everyone finds desirable. We call such information "consensus solutions." Consensus solutions are useful in shared environments, such as broadcast music or artistic displays in public spaces. We will outline the relationship between a proposed IIS and a natural immune system, propose some applications of an IIS, including information filtering, interactive design, and collaborative design, then summarize the results of an experiment testing an IIS implementation.

2 Related work

Several areas of research have influenced our conception of an IIS. An IIS must be able to learn from past encounters, and the issues of learning and memory have long been addressed by the fields of artificial intelligence and machine learning. The primary task that we propose for an IIS, information filtering, has been explored by the field of human-computer interaction. Collaborative filtering may be relevant for IISs that classify data that are difficult to evaluate algorithmically. A few collaborative filtering systems make recommendations to groups instead of individuals. These group recommender systems perform a function similar to a group IIS. Finally, an IIS should be informed by

earlier work in artificial immune systems. All of these influences are briefly discussed below.

Case-based reasoning is a technique that adapts solutions to past problems to solve similar current problems [44]. Memory-based reasoning [49] and instance-based learning [1] are related schemes that use the solution of the most similar previous problem. Systems using these approaches learn by "remembering" specific past events rather than creating rules or generalizations. Immune memory uses a form of instance-based learning; the particular response that was effective in clearing a pathogen will likely be used in future encounters with related pathogens [29, 43, 5].

Associative memories, often called content-addressable memories, are neurally inspired architectures that can retrieve items using approximate addresses. Smith outlines the parallels between Kanerva's sparse distributed memory [25] and the memory of the natural immune system [47]. The memory of the natural immune system is not exact, and exposure to a novel pathogen can elicit the response primed by a related pathogen.

The term "information filtering" refers to a large range of techniques used to remove data from an incoming stream on the basis of user- or group- specified preferences [2]. Early approaches used simple rules [30] or signatures (e.g. keywords) to identify undesirable data to block. These approaches are still popular, and many commercial products, such as Cyberpatrol [9] for web content and the Realtime Blackhole List [41] and Brightmail [6] for e-mail, come with long lists of rules and signatures, which can be effective in blocking undesirable data but are vulnerable to malicious sources that can craft information to bypass them. To thwart these adaptive adversaries and to personalize the filtering, the user is often allowed to specify additional rules for accepting and rejecting data. Unfortunately, the specification of such rules is often difficult and error-prone, and therefore not used routinely. An IIS should incorporate reliable signatures of undesirable data as a first line of defense to be supplemented with more adaptive techniques to provide better and more personalized coverage.

Several research systems simplify the filter specification problem by placing the burden of generating rules on software rather than on a user or programmer. Infoscope [14] monitors a user's behavior to create rules for Usenet newsgroup filtering. The system suggests these rules to the user, who can accept, modify, or reject them. Maxims [34], an interface agent for e-mail, also generates filtering rules based on user behavior, but it suggests actions for the user to take rather than rules when it is confident in its predictions. Rule-based learning schemes often require many examples before they can infer new rules. In contrast, an IIS using an instance-based learning approach could learn to block a class of data upon seeing only a single exemplar.

Collaborative filtering uses the preferences of others to help an individual make choices [31, 16, 42]. For example, a collaborative filtering system would recommend an item for a person to purchase by choosing an item purchased by someone with a similar purchase history. By harnessing the collective preferences of many individuals, such systems can infer similarity between items without needing to understand the relationship between them. This approach is useful when it is difficult for a program to determine similarities between items, such as art or music. An IIS could incorporate collaborative filtering techniques to determine the similarity between items for its associative memory capabilities.

There are a few systems that recommend items to groups instead of individuals. MusicFX [32] selects music stations that are broadcast to a gym full of people. The members of the gym must rate all the stations beforehand, and MusicFX plays one of the stations with the highest average rating. One shortcoming of MusicFX is that it apparently does not scale to a large number of choices. If the users were not able to evaluate all of the stations, the quality of the system's choices would likely be degraded. GroupCast [33], developed by the same research group, used a conceptually similar scheme to display content on a public display system. Unfortunately, they found that the necessary user profiles would have been too large for any user with a reasonable amount of patience to complete. In addition, without extensive profiles it was difficult to find appropriate intersections of user preferences to put on the GroupCast displays. Instead, they displayed content that was interesting to *one* of the users, hoping that by chance others would have similar interests. PolyLens [36] recommends movies to small groups of people who watch movies together. This system applies a standard collaborative filtering algorithm to make recommendations for each of the group members then combines the results to make a group recommendation. These systems give insight into the nature of finding solutions for groups. Notably, it is difficult to make recommendations that satisfy all members of a large group.

Immune system inspired algorithms have often been used for anomaly detection. They draw on the metaphor of the adaptive immune system's ability to distinguish between *self*, or normal data, and *nonself*, or anomalous data. One of the first such systems was the negative selection algorithm introduced by Forrest *et al* [15]. The algorithm generated random strings and those that were similar to sequences of bytes in a given computer file were eliminated. The surviving strings were therefore not similar to any in the file. If one of these strings ever matched the contents of the file, then this indicated that the contents had been changed since the training period. These strings were used as *negative detectors* to detect novel sequences of bytes, such as those introduced when a virus corrupts or infects a file. The ARTIS framework is an extension of this work

that applies negative selection to detect anomalies in streams of data rather than in static data sets [18, 19]. This framework was used to create systems to detect network intrusions [18, 19, 27, 54].

We believe that most useful sources of information present continually changing streams of data, so that it would be undesirable for an IIS to reject all novel data. The IIS is inspired by the immune system's ability to remember past encounters with pathogens, while the artificial immune system approach to anomaly detection is usually based on the immune system's ability to detect novel foreign proteins. The anomaly detection ability of ARTIS could complement an IIS for certain applications, but for many applications we imagine using negative detectors without negative selection.

Many computer scientists have developed artificial immune systems based on idiotypic network theory [24]. The idiotypic systems focus on the dynamics of the interactions among similar antibodies and antigens. Although many do not attempt to reproduce the behaviors seen in the natural immune system, they have useful properties that have been applied to search [4], data classification [21], cluster detection [50], and data mining [12]. The classifications produced by idiotypic artificial immune systems could potentially be used as metadata to enhance the discrimination of an IIS.

3 The immune system as an information filter

We believe that an IIS can borrow several pattern recognition mechanisms from the natural immune system. Our natural immune system consists of two components that use different pathogen recognition strategies. The *innate* immune system uses a few reliable signatures of foreignness to identify invaders, which Janeway calls pathogen-associated molecular patterns (PAMP) [22]. An example of a PAMP is the mannose carbohydrate molecules found on many bacteria and other pathogens but not in mammals [48]. These signatures have been stable over evolutionary time and are encoded in the genome of our immune systems. This strategy is used by many of the signature and rule based information filtering products mentioned in Section 2. These products could serve as a first line of defense, playing the role of the innate immune system in an IIS. However, not all signatures of pathogens have been (or even can be) anticipated, and evolution will favor pathogens that do not carry the signatures recognized by our innate immune systems. One role of the *adaptive* immune system, discussed below and outlined in Table 1, is to discover the signatures of pathogens not covered by the innate immune system. In the following subsections we describe some issues that an IIS must face and how one can draw inspiration from the natural immune system to address them.

Natural IS	Information IS
shape space	parameter space
self	desirable information
non-self	undesirable information
helper T cell	user's judgment
costimulation	rejection of information by user
naïve cells	implicit (not instantiated)
active lymphocyte	detector
memory lymphocyte	detector
cytolytic activity	censor solution
cross-reactive radius	detector radius
thymic selection	protecting known desirable information
illness	user exposed to undesirable data

Table 1: The immunological analogy made explicit.

3.1 Negative detectors and shape space

An IIS should be able to remember which pieces of information a user rejected in the past so it can censor them in the future. However, the strategy of rejecting each item individually is ineffective when one is faced with a seemingly limitless variety of information. An IIS must be able to generalize; rejecting one item should implicitly reject similar items. The natural immune system has this ability.

The adaptive immune system has a repertoire of lymphocytes that detect pathogens. Each lymphocyte is specific to a particular antigen, or protein signature, expressed by pathogens. If a lymphocyte detects a cell with a matching signature, it may destroy it. However, pathogens may mutate and subtly change their antigenic profiles, so lymphocytes should also be able to recognize close variants. Perelson and Oster suggested the conceptual framework of *shape space* [38], a high-dimensional space that represents the universe of possible antigens. Every antigen has a location in shape space, and small mutations in a pathogen may alter its proteins, thus shifting its location in shape space. For a lymphocyte to be effective, it should be able to cover a large enough area in shape space that most mutations would not evade detection. The area in shape space that a lymphocyte covers is sometimes known as its *ball of stimulation* because it is postulated that a lymphocyte can recognize an antigen within a certain radius of its location in shape space.

An IIS could use negative detectors to censor information that the user does not want. As with the natural immune system, a detector should be able to cover a volume in shape space, not just a point. Therefore, it is necessary for an IIS to have some notion of the similarity between two pieces of information. Two items that are similar are close in "information space." Collaborative filtering techniques could be used in cases in which it is too difficult to define a function that en-

codes the subjective similarity between two pieces of information.

3.2 Costimulation

Because everyone has different informational needs, each IIS user should be able to decide which types of data to reject. Many information filters require the user to write rules to customize the filtering, but we believe that the user should need only to identify exemplars of undesirable information. Once the user rejects a piece of information, an IIS should be able to automatically reject similar information in the future.

In the adaptive immune system, helper T cells are generally required to *costimulate*, or activate, cells in the presence of a novel pathogen. Helper T cells provide confirmation that a pathogen should be eliminated. This process reduces the chances of immune cells attacking the body, which is known as an autoimmune response. Once costimulated, the effector cell becomes active and can attempt to eliminate the invader, whether by releasing antibodies in the case of B cells or by killing the infected cells directly in the case of cytotoxic T cells. Some co-stimulated cells become memory cells, which are long-lived. In future encounters with the same pathogen, memory cells have lesser or even no costimulation requirement.

In an IIS, the user could adopt the role of the helper T cells by providing costimulation signals to the system, an idea introduced in [19]. The idle cells waiting for costimulation are implicit—only detectors corresponding to active or memory cells need to be instantiated. When the user rejects a piece of information, a detector specific to that item would be created. These detectors would prevent any similar data from being presented in the future. The user's only responsibility would be to inform the IIS when undesirable data are being presented.

3.3 The addition of negative selection

When the user has rejected a sufficient amount of information, the space not covered by detectors approximates the space of useful information (Figure 1). Unfortunately, useful information that is too similar to unwanted information runs the risk of being censored by an IIS negative detector. Therefore, we suggest incorporating a technique that the adaptive immune system uses to prevent the immune system from attacking the body's own cells.

The adaptive immune system uses thymic selection to eliminate T cells that may harm the body. Before T cells can enter the repertoire, they are exposed to a large sample of the body's own proteins. Those that bind too tightly to one of the body's proteins are eliminated in a process known as negative selection. Therefore, the T cells that survive are not likely to recognize a self protein.

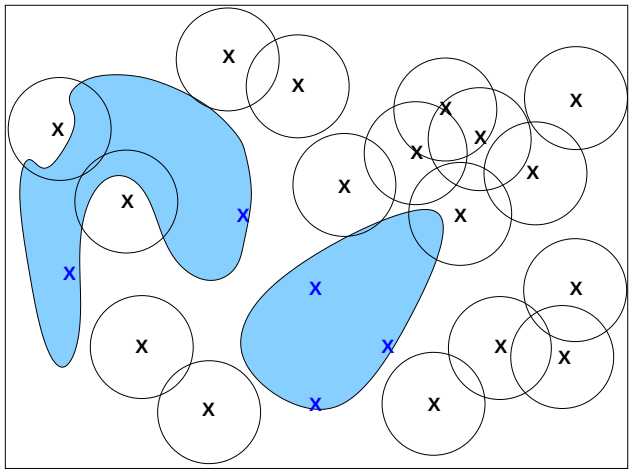

Figure 1: Coverage in the model without negative selection. The shaded regions represent information that is useful. The circles represent the extent of active detector coverage. The Xs without circles represent the detectors that should never be costimulated because they are within the regions of acceptable solutions.

A similar strategy could be employed in advance by an IIS to protect types of information known to be useful. These types could be declared "off-limits" to the IIS and would be allowed to bypass the IIS to reach the user. This is especially useful when the characteristics of certain desirable information are known *a priori*. For example, the IISs of a company's employees should probably not be allowed to eliminate official company e-mail. When a user costimulates a solution whose detector would cover some desirable information, the system could ignore the costimulation signal because there should be no "implicit" detectors in this region. No information from the "good" regions of information space will ever be censored by the detectors (Figure 2).

3.4 The role of senescence

Users may want to filter out some types of information for only a short period of time. For example, if a radio station plays a song too frequently or if a news story receives too much coverage, a listener may tire of it. These individuals may actually enjoy hearing the song or listening to new developments in the news story at a later date, so the detectors would be counterproductive after their "natural" lifetimes.

Active immune cells have short lifetimes, and memory cells can be eliminated by competition for space [45]. These features may be desirable in the algorithm for two reasons. The first is to provide "rolling coverage" of self. If the fitness function (e.g. the user's tastes) change over time, one could have the lifetime of the active immune cells be finite to reflect the dynamic nature of the user's judgment. The second reason is space efficiency. It may not be feasible to store an unbounded number of detectors. One could "age out" old detec-

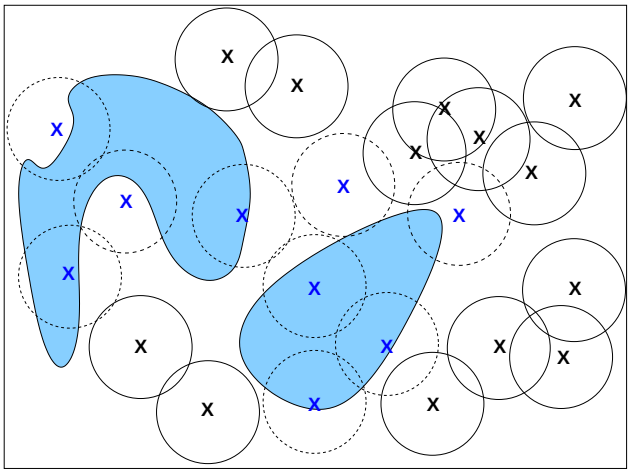

Figure 2: Coverage in the model using negative selection. The dotted circles represent the extent of detectors that are eliminated by negative selection. The solid circles are regions covered by active detectors. Note that none of the useful information protected by negative selection (the shaded regions) can be covered by detectors.

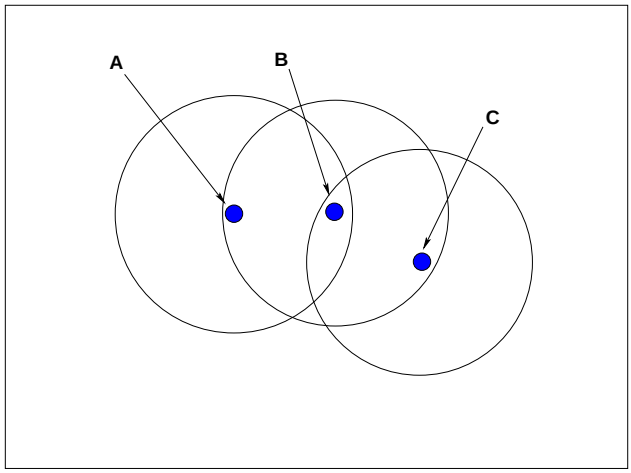

Figure 3: The effect of history. The dots labeled "A" and "B" and "C" represent solutions the user does not like. The circles are the extents of their negative detectors, or "balls of stimulation". If solution A is rejected by the user first, the detector that forms around it would reject B before it could be presented. However, C could be presented because it does not fall within the scope of the detector for A. However, if B had been presented first, the story would be different. Neither A nor C would be seen after B because its detector would cover all three solutions.

tors to make room for new ones. Alternatively, the user could manually create memory detectors to cover patterns that he or she never wants to see again.

3.5 The effect of history

The order in which an IIS is exposed to information can have impact on its effectiveness. Such phenomena have been observed in the natural immune system, particularly in the case of influenza. Immunologists have discovered that the response to a strain of flu may be dominated by cells that were created in response to an earlier exposure to a *different* strain [10, 13]. These memory cells are probably most effective against the strain that generated them, but they can respond to related ones. This phenomenon is known as original antigenic sin, and many vaccines take advantage of this effect. For example, if one is exposed to the relatively harmless cowpox bacteria, one is protected against the related but deadly smallpox [23]. Unfortunately, prior exposure to antigens can also work against us [46]. For example, a flu vaccine works by eliciting a mild response to a particular strain's flu antigens so that an individual will be able to mount an effective secondary response when exposed to it in the future. However, the memory cells created by a vaccine from a previous year may attack and eliminate subsequent vaccines before they can establish protective immunity. If the first vaccine does not provide protection against the strains corresponding to these later vaccinations, this individual would be vulnerable to them (Figure 3). If this individual had not received this first vaccine, the subsequent vaccines could have been effective.

One should be able to "vaccinate" an IIS by exposing it to undesirable information without necessarily exposing the user. This would allow an administrator to preemptively block the passage of certain kinds of information to a user. For example, a corporation might prohibit certain kinds of e-mail or web traffic, such as pornography or personal e-mail. The corporation could "vaccinate" the IISs of its employees with exemplars from the banned categories, and the employees would not be exposed to these kinds of information. Because the order in which an individual is exposed to undesirable information may affect the coverage of the individual's IIS, the vaccination strategy should be planned with care.

4 Applications

The most obvious use of an IIS of the sort described here is information filtering. An IIS could serve as a personalized interface agent that learns a user's preferences for sources of information or for a range of options that is too large or dynamic for a user to evaluate. Because it only requires feedback when the user is exposed to something he or she does not want and it learns without using separated training and testing phases, an IIS could be a non-intrusive addition to many user interfaces. It could complement active strategies, such as information retrieval, that search for potentially useful information.

The IISs of individuals can be combined to produce a group IIS. One can think of an IIS as a sieve that fil-

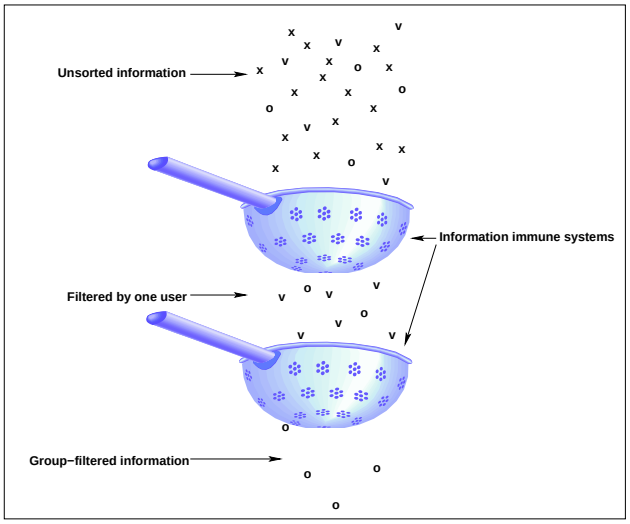

Figure 4: An information immune system as a sieve. The IIS stands between a stream of information and a user, blocking a significant portion of it. Only the information that can pass through the "sieve" actually reaches the user. When the IISs of multiple users are applied in serial, the information that passes all IISs are "consensus solutions".

ters undesirable information. The data that can pass through the "sieves" of many people are those that are likely to satisfy all of them (Figure 4). We call these data *consensus solutions*. A group IIS would be useful when the group is exposed to shared information. For example, if co-located people want to listen to music together, one would want to play music that none of the individuals dislikes. It remains to be seen how well a group IIS will work with a large group of users in a particular domain. Consensus-finding will become more important with the increase in the number of intelligent environments that automatically respond to the users' needs. For example, smart home technology can adjust the music, artwork, temperature, and lighting to accommodate its occupants. Most research focuses on catering to a single individual [26, 17, 51, 20], but for many environments it will be important to satisfy the preferences of multiple occupants.

An IIS could be used to assist designers and artists [7]. If a random source of design solutions or works of art were fed to an IIS, only those that are not similar to those rejected in the past would pass through. The quality of the solutions from this filtered stream should be significantly better than the unfiltered stream. This could be a useful strategy for design problems in which a designer or artist is interested in exploring a large range of possible solutions. The solutions could be refined or optimized using other techniques, such as evolutionary design [3].

Collaborative design could be facilitated by using the IISs of multiple individuals. The combination of IISs is the superset of solutions that people dislike. Consensus solutions are not optimal solutions, but a variety of solutions that are "good enough" for *everyone*. In certain cases, it would be preferable to combine the favored solutions of each of the group members instead of using a group IIS. This could be done by taking the intersection of the favored solutions of the members or by combining (hybridizing) them. The former strategy is problematic when the solution space is too large for a user to specify the set of all acceptable solutions (as was found with GroupCast[33]) or if the knowledge of a user's preferences is incomplete. In these cases, intersections will be difficult to find. The latter strategy of combining solutions can be difficult. It is often not obvious how to combine the desirable traits of two solutions to produce a third good solution. By combining the *dislikes* of multiple users, the space of potential candidate solutions is likely to be larger and there is no need to combine solutions.

5 An example: An aesthetic information immune system

We have applied the principles discussed in this paper to design a simple IIS that generates computer art [7], and we summarize the results here. The IIS characterized several users' preferences for a particular family of computer-generated images known as Biomorphs [11]. Biomorphs are recursively drawn figures that can be defined by nine parameters. Each user was shown a set of randomly generated Biomorphs and instructed to reject those that he or she did not like. For each user, an IIS was created based on the parameters of the rejected Biomorphs. The IIS filtered out any images that had parameters similar to those rejected in the past, and they formed a rough estimate of the parts of Biomorph parameter space that each user wanted to avoid.

We tested whether a user could use an IIS to filter a stream of randomly generated Biomorphs to produce an edited stream of high quality Biomorphs, based on the subjective judgments of the user. We also investigated group IISs that applied the IISs of several users in serial. We wanted to determine if the addition of other users' IISs would enhance or degrade the quality of a single user's IIS. These effects were measured by having the users evaluate three sets of randomly generated Biomorphs that were filtered using no IIS, their own IIS, a group IIS composed of seven users' IISs. Most users preferred the Biomorph images filtered using their own IISs to the unfiltered ones, suggesting that the IISs had preferentially filtered out images that would have been rejected by the users. The group IIS was less successful, possibly because of differences among the users' Biomorph aesthetic preferences or possibly because of the coarseness of the detectors (we used quite coarse-grained detectors in order to reduce the training time for each user). We repeated the test with a subset of three users and a group IIS

composed of only these three users' IISs. The images produced by this smaller group's IIS were perceived to be better than unfiltered, and each user found these images to be no worse than those produced using their own IISs, indicating the possibility of a consensus solution.

6 Conclusions

We believe that information immune systems could play an important role in this age of information overload. To date, we as a society have developed only crude coping mechanisms to allow us to survive the enormous amounts of data to which we are routinely exposed [35]. A successful IIS would reduce the load and make other strategies for finding and processing information more effective.

Information immune systems, however, should be fielded with caution. As filtering strategies become more sophisticated, the producers of unwanted information will themselves adapt, creating a kind of information arms race. We see this already in the adaptation of magazine advertisements designed to resemble content articles and "junk mail" packaged in official-looking envelopes. Even more insidious techniques embed advertising in content in which people are interested. Advertisements can be wrapped around e-mail for presentation before the user can receive it [8], corporate logos and products can be digitally edited into films and television programs [53], and some shows integrate their sponsors' products into the plotlines [37]. Even in the absence of adaptive adversaries, our information filtering technology will drive a selective process that will minimize the differences between desirable and undesirable information. As our filters gain efficacy, undesirable information will evolve to evade them. The filters must constantly co-evolve or else they will rapidly become useless. When we begin deploying IISs, we must be prepared to live in a dynamic information ecosystem in which our defenses must adapt as quickly as the abilities of unwanted information to penetrate them [52].

Acknowledgments

We thank Marc Millier of Intel Corporation for suggesting the term "information immune system." The authors gratefully acknowledge the support of the National Science Foundation (ANIR-9986555), the Office of Naval Research (N00014-99-1-0417), Defense Advanced Projects Agency (AGR F30602-00-2-0584), the Intel Corporation, and the Santa Fe Institute.

References

[1] D. Aha, D. W. Kibler, and M. K. Albert. Instance-based learning algorithms. *Machine Learning*, 6:37–66, 1991.

[2] N. J. Belkin and W. B. Croft. Information filtering and information retrieval: Two sides of the same coin? *Communications of the ACM*, 35(12):29–38, 1992.

[3] P. J. Bentley, editor. *Evolutionary Design by Computers*. Morgan Kaufmann Publishers, San Francisco, California, 1999.

[4] H. Bersini and F. J. Varela. Hints for adaptive problem solving gleaned from immune networks. In H. Schwefel and R. Männer, editors, *Parallel Problem Solving from Nature*, pages 343–354, Berlin, 1991. Springer-Verlag.

[5] J. A. Borghans, A. J. Noest, and R. J. De Boer. How specific should immunological memory be? *J Immunol*, 163(2):569–75, 1999.

[6] *Brightmail Solution Suite*. Brightmail, Inc., San Francisco, California, 2002. http://www.brightmail.com.

[7] D. L. Chao and S. Forrest. Generating biomorphs with an aesthetic immune system. In *Artificial Life VIII: Proceedings of the Eighth International Conference on the Simulation and Synthesis of Living Systems*, 2002. (in press).

[8] N. Cochrane. Mobile entrepreneur rapt in wireless e-mail advertising. *The Age (Melbourne)*, 24 Jul 2001:7.

[9] *Cyberpatrol*. SurfControl plc, Westborough, Massachusetts, 2002.

[10] F. M. Davenport, A. V. Hennessy, and T. Francis. Epidemiologic and immunologic significance of age distribution to antibody to antigenic variants of influenza virus. *J Exp Med*, 98:641–656, 1953.

[11] R. Dawkins. *The Blind Watchmaker*. Longman Scientific and Technical, Harlow, UK, 1986.

[12] L. N. De Castro and F. J. Von Zuben. aiNet: An artificial immune network for data analysis. In H. A. Abbass, R. A. Sarker, and C. S. Newton, editors, *Data Mining: A heuristic approach*, chapter XII, pages 231–259. Idea Group Publishing, USA, Hershey, Pennsylvania, 2001.

[13] S. Fazekas de St. Groth and R. G. Webster. Disquisitions of original antigenic sin. I. Evidence in man. *J Exp Med*, 124(3):331–45, 1966.

[14] G. Fischer and C. Stevens. Information access in complex, poorly structured information spaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 1991)*, pages 63–70, 1991.

[15] S. Forrest, A. S. Perelson, L. Allen, and R. Cherukuri. Self-nonself discrimination in a computer. In *Proceedings of the 1994 IEEE*

Symposium on Research in Security and Privacy. IEEE Computer Society Press, 1994.

[16] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–70, 1992.

[17] S. R. Hedberg. After desktop computing: A progress report on smart environments research. *IEEE Intelligent Systems*, 15(5):7–9, 2000.

[18] S. A. Hofmeyr. *An immunological model of distributed detection and its application to computer security*. PhD thesis, University of New Mexico, Albuquerque, New Mexico, 1999.

[19] S. A. Hofmeyr and S. Forrest. Architecture for an artificial immune system. *Evolutionary Computation*, 7(1):1289–1296, 1999.

[20] *House_n Living Laboratory*. School of Architecture and Planning, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2001.

[21] J. E. Hunt and D. E. Cooke. Learning using an artificial immune system. *Journal of Network and Computer Applications*, 19:189–212, 1996.

[22] C. A. Janeway Jr. The immune system evolved to discriminate infectious nonself from noninfectious self. *Immunol Today*, 13(1):11–6, 1992.

[23] E. Jenner. *An Inquiry into the Causes and Effects of the Variolae Vaccinae; a Disease Discovered in some of the Western Counties of England, Particularly Gloucestershire, and Known by the Name of the Cow Pox*. 1798.

[24] N. K. Jerne. Towards a network theory of the immune system. *Ann Immunol (Inst Pasteur)*, 125C:373–389, 1974.

[25] P. Kanerva. *Sparse Distributed Memory*. MIT Press, 1988.

[26] C. D. Kidd, R. J. Orr, G. D. Abowd, C. G. Atkeson, I. A. Essa, B. MacIntyre, E. Mynatt, T. E. Starner, and W. Newstetter. The aware home: A living laboratory for ubiquitous computing research. In *Proceedings of the Second International Workshop on Cooperative Buildings (CoBuild'99)*, pages 191–198, 1999.

[27] J. Kim and P. J. Bentley. The artificial immune model for network intrusion detection. In *Proceedings of the 7th European Conference on Intelligent Techniques and Soft Computing (EUFIT'99)*, 1999.

[28] P. Lyman and H. R. Varian. How much information?, 2000. Retrieved from http://www.sims.berkeley.edu/how-much-info.

[29] C. R. Mackay, W. L. Marston, L. Dudler, O. Spertini, T. F. Tedder, and W. R. Hein. Tissue-specific migration pathways by phenotypically distinct subpopulations of memory T cells. *Eur J Immunol*, 22(4):887–95, 1992.

[30] T. W. Malone, K. R. Grant, and F. A. Turbak. The Information Lens: An intelligent system for information sharing in organizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 1986)*, pages 1–8, 1986.

[31] T. W. Malone, K. R. Grant, F. A. Turbak, S. A. Brobst, and M. D. Cohen. Intelligent information sharing systems. *Communications of the ACM*, 30(5):390–402, 1987.

[32] J. F. McCarthy and T. D. Anagnost. MusicFX: An arbiter of group preferences for computer supported collaborative workouts. In *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work*, pages 363–372. ACM Press, 1998.

[33] J. F. McCarthy, T. J. Costa, and E. S. Liongosari. UniCast, OutCast & GroupCast: An exploration of new interaction paradigms for ubiquitous, peripheral displays. In *Workshop on Distributed and Disappearing User Interfaces in Ubiquitous Computing at the SIGCHI Conference on Human Factors in Computer Systems (CHI 2001)*. ACM Press, 2001.

[34] M. Metral. *A Generic Learning Interface Agent*. B.Sc. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1992.

[35] J. G. Miller. Information input overload and psychopathology. *American Journal of Psychiatry*, 116(8):695–704, 1960.

[36] M. O'Connor, D. Cosley, J. A. Konstan, and J. Riedl. PolyLens: A recommender system for groups of users. In *Proceedings of the 7th European Conference on Computer Supported Cooperative Work (ECSCW 2001)*, pages 199–218. Kluwer Academic, 2001.

[37] G. Pennington. Just try zapping these ads. *St. Louis Post-Dispatch*, 14 Apr 2002:F1.

[38] A. S. Perelson and G. F. Oster. Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *J Theor Biol*, 81(4):645–70, 1979.

[39] N. Postman. Informing ourselves to death, 11 Oct 90. Speech delivered to the German Informatics Society (Gesellschaft für Informatik).

[40] N. Postman. *Technopoly. The Surrender of Culture to Technology.* Vintage Books, New York, 1992.

[41] *Realtime Blackhole List.* Mail Abuse Prevention System LLC, Redwood City, California, 2002. http://www.mail-abuse.org/rbl/.

[42] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl. GroupLens: An open architecture for collaborative filtering of Netnews. In *Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work*, pages 175–186. ACM Press, 1994.

[43] F. Sallusto, D. Lenig, R. Forster, M. Lipp, and A. Lanzavecchia. Two subsets of memory T lymphocytes with distinct homing potentials and effector functions. *Nature*, 401(6754):708–12, 1999.

[44] R. C. Schank. *Dynamic Memory: A theory of reminding and learning in computers and people.* Cambridge University Press, New York, 1982.

[45] L. K. Selin, K. Vergilis, R. M. Welsh, and S. R. Nahill. Reduction of otherwise remarkably stable virus-specific cytotoxic T lymphocyte memory by heterologous viral infections. *J Exp Med*, 183(6):2489–99, 1996.

[46] D. J. Smith, S. Forrest, D. H. Ackley, and A. S. Perelson. Variable efficacy of repeated annual influenza vaccination. *Proc Natl Acad Sci U S A*, 96(24):14001–6, 1999.

[47] D. J. Smith, S. Forrest, and A. S. Perelson. Immunological memory is associative. In *Workshop Notes, Workshop 4: Immunity Based Systems, Intnl. Conf. on Multiagent Systems*, pages 62–70, 1998.

[48] P. D. Stahl and R. A. Ezekowitz. The mannose receptor is a pattern recognition receptor involved in host defense. *Curr Opin Immunol*, 10(1):50–5, 1998.

[49] C. Stanfill and D. Waltz. Toward memory-based reasoning. *Communications of the ACM*, 29(12):1213–1228, 1986.

[50] J. Timmis. *Artificial immune systems: A novel data analysis technique inspired by the immune network theory.* PhD thesis, University of Wales, 2000.

[51] The user in control. *Philips Research Password*, 3:10–13, 2000.

[52] L. Van Valen. A new evolutionary law. *Evolutionary Theory*, 1:1–30, 1973.

[53] Virtual ads, real problems. *Advertising Age*, 70(22):30, 1999.

[54] P. D. Williams, K. P. Anchor, J. L. Bebo, G. H. Gunsch, and G. D. Lamont. CDIS: Towards a computer immune system for detecting network intrusions. In W. Lee, L. Me, and A. Wespi, editors, *4th International Symposium, Recent Advances in Intrusion Detection*, pages 117–133, Berlin, 2001. Springer-Verlag.